

WRF EMS Installation and benchmark case run

For this exercise, you will install the WRF EMS on your workstation and then run the benchmark case for either the NMM or ARW core. The exercise is straight forward with step-by-step instructions provided below. Feel free to ask for assistance should you have any problems.

I. WRF EMS installation

The WRF EMS includes an installation script, `wrf_install.pl`, which is used to simplify the process of installing and configuring the package on your system. It is recommended that the installation be conducted by user `root`; however, it may be installed by another user provided that you have full write permission on the partition where the package will reside.

*For the purpose of this exercise you will install the WRF EMS as **root** user.*

To install the WRF EMS complete the following steps (as **root**):

Do: Mount the WRF EMS DVD on your workstation

E.g.: `# mount /mnt/cdrom`

Do: Change directories to your DVD drive

E.g.: `# cd /mnt/cdrom`

Note: The drive should be mounted with executable permissions; otherwise you will have problems when running the wrf install routine.

Do: Run the "`wrf_install.pl`" script

E.g.: # ./wrf_install.pl

During the installation process the wrf_install.pl routine will attempt to do the following:

- Prompt you for the installation directory (**Default: /usr1**)
- Check to make sure the directory exists and whether you have write permission
- Prompt you for the name of the user to assign ownership of the package
- Create a new account and home directory if the user does not exist
- Check that the user's login shell is either tcsh or csh
- Prompt you for a password if a new user was created
- Install the WRF EMS from the DVD

Note: Should the WRF EMS DVD be mounted without executable permission, then you will get an error when running the "wrf_install.pl" routine. An easy solution to this problem is to copy the wrf_install.pl script from the DVD to the hard drive and run the routine from that location. You will be prompted as to the location of the WRF EMS package at which time simply specify the mounted DVD.

While the installation proceeds, useful information will be printed to the screen so don't leave the room.

Following installation of the files, the script will attempt to complete the following:

- Build PerlTk libraries for your system. Should the PerlTk build fail, precompiled libraries for either 32- or 64-bit system will be installed.
- Configure the WRF EMS, which includes updating configuration files and Perl scripts with the location of the WRF EMS on your system.
- Attempt to configure SSH on your system so that you can run distributed memory (DM) binaries. These binaries will provide a 15% performance gain over the shared memory (SM) executables with the ARW and NMM cores.

- Install the wrf_update.pl routine in the user's crontab file (disabled). The wrf_update.pl script automatically downloads and installs updates and patches to the EMS from the SOO/STRC ftp server and notifies the user(s) of any changes.
- Update the user's .cshrc file so that the WRF EMS environment variables are correctly defined the next time the user logs in.

See, isn't that simple? The WRF EMS practically installs itself.

Depending on the speed of your DVD reader and hard disks, the installation should take ~10 to 15 minutes. Most of this time is spent unpacking the very large surface data sets such as topography and land use.

Following the successful installation, log out and return as the WRF EMS user.

Do: log out as user rot and log in as wrf user

To make sure your environment is set correctly, try the following commands:

Do: cd \$WRF - you should be in the top level of the WRF EMS

Do: cd \$WRF_UTIL - You will see the contents of the \$WRF/util directory

Provided that the above commands worked you are ready to proceed to the second part of the exercise.

II. Running a WRF EMS benchmark case

Beneath the wrf/util directory you will see two subdirectories each containing a preconfigured case for the NMM and ARW cores for the purpose of benchmarking your system. The **bench_nmm** and **bench_arw** directories contain the configuration for the

NMM and ARW cores respectively. The case for both cores is the same with the only difference in the computational domains being the horizontal grid point dimensions specified.

The benchmark case is a 24-hour simulation of oceanic cyclogenesis event that occurred over the Northeast US (centered over Boston, MA) from January 2005. Both the NMM and ARW core benchmarks were designed to run over 5005 horizontal grid points (NMM:55x91, ARW:77x65) and 45 vertical levels at 15km grid spacing.

For the purpose of this exercise you will choose either the ARW or NMM core to benchmark, and a memory environment (Shared, SM or distributed, DM). Following completion of the simulation you will compare the time required to complete the run with others. ***Fun, fun, fun!***

Do: Decide which core (ARW or NMM) you will run

Do: Change to the appropriate directory used for running the benchmark you selected

Do: Decide which memory environment you will use. Your choices for the memory environment are Shared (SM) or Distributed (DM).

If you choose distributed memory (DM):

You will have to make sure your system is configured correctly to use passwordless ssh. A simple test is to run the following command:

% ssh `hostname` date

If the command returns the date without being prompted for a password you are set. If you are prompted with an authorization message then simply respond “yes”, enter the password and then try the above command again.

If you are prompted for a password each time you run the above command then you will be unable to use distributed memory until passwordless ssh is correctly configured on your system.

If the above command was successful then you will need to decide the number of processors on which to run, which is controlled by the MPIRUNARGS parameter in the conf/wrf_run/run_DMrun.conf file:

MPIRUNARGS = "-arch LINUX -np 8 -all-local"

For this exercise, the “-np” value should default to “8”, i.e, “-np 8”, which is the number of physical CPUs (2) multiplied by the number of cores per CPU (4). You may change the value from the default; however, then make sure you provide this information to the class during the discussion section because it will impact your results.

If you choose shared memory (SM):

The number of CPUs that you will use when running the simulation is governed by the NCPUS environment variable. To see the current value, running

% echo \$NCPUS

should return a value of “2”, which is the **number of physical CPUs** on your system. The value of NCPUS is set during the configuration part of the installation and is contained in the wrf/WRF.cshrc file:

setenv NCPUS 2

If you wish to change this value then edit the WRF.cshrc file and either run the “source” command on the file or simply log out and back in again. Check to see if the change took effect with the “echo \$NCPUS” command.

Running the benchmark case

Running the benchmark case is simple:

Step 1 *From the benchmark directory run wrf_prep*

% wrf_prep --benchmark

Step 2a Run the model (Shared memory)

% wrf_run –benchmark --SM

Note that the default is to use the shared memory executable (--SM option) so you do not need the “–SM”

Step 2b Run the model (Distributed memory)

% wrf_run –benchmark --DM

The amount of time required for the benchmark case to complete will depend upon the core used and the memory environment.

Upon completion of the run there will be a file in the static directory called benchmark.results that contains your results.

Feel free to send this file to the SOO Science and Training Resource Coordinator for inclusion in the SOO/STRC WRF benchmarks site:

http://strc.comet.ucar.edu/wrf/wrf_benchmarking.htm

III. Compare the results of your run with those of the other experiments

Briefly discuss with the other groups the results of the benchmark runs

End of Exercise #1